

Bridging the Training-Inference Gap for Dense Phrase Retrieval

Gyuwan Kim¹ Jinhyuk Lee^{2*} Barlas Oğuz³ Wenhan Xiong³

Yizhe Zhang^{3†} Yashar Mehdad³ William Yang Wang¹

¹University of California, Santa Barbara ²Korea University ³Meta AI
gyuwankim@ucsb.edu, jinhyuk_lee@korea.ac.kr
{barlaso, xwhan, yizhezhang, mehdad}@fb.com, william@cs.ucsb.edu

Abstract

Building dense retrievers requires a series of standard procedures, including training and validating neural models and creating indexes for efficient search. However, these procedures are often misaligned in that training objectives do not exactly reflect the retrieval scenario at inference time. In this paper, we explore how the gap between training and inference in dense retrieval can be reduced, focusing on dense phrase retrieval (Lee et al., 2021a) where billions of representations are indexed at inference. Since validating every dense retriever with a large-scale index is practically infeasible, we propose an efficient way of validating dense retrievers using a small subset of the entire corpus. This allows us to validate various training strategies, including unifying contrastive loss terms and using hard negatives for phrase retrieval, which largely reduces the training-inference discrepancy. As a result, we improve top-1 phrase retrieval accuracy by 2~3 points and top-20 passage retrieval accuracy by 2~4 points for open-domain question answering. Our work urges modeling dense retrievers with careful consideration of training and inference via efficient validation while advancing phrase retrieval as a general solution for dense retrieval.

1 Introduction

Dense retrieval aims to learn effective representations of queries and documents by making representations of relevant query-document pairs to be similar (Chopra et al., 2005; Van den Oord et al., 2018). With the success of dense passage retrieval for open-domain question answering (QA) (Lee et al., 2019; Karpukhin et al., 2020), recent studies build an index for a finer granularity such as dense *phrase* retrieval (Lee et al., 2021a), which largely improves the computational efficiency of open-domain QA

by replacing the retriever-reader model (Chen et al., 2017) with a retriever-only model (Seo et al., 2019; Lewis et al., 2021). Also, phrase retrieval provides a unifying solution for multi-granularity retrieval ranging from open-domain QA (formulated as retrieving phrases) to document retrieval (Lee et al., 2021b), which makes it particularly attractive.

Building a dense retrieval system involves multiple steps (Figure 1) including training a dual encoder (§4), selecting the best model with validation (§3), and constructing an index (often with filtering) for an efficient search (§5). However, these components are somewhat loosely connected to each other. For example, model training is not directly optimizing the retrieval performance using the full corpus on which models should be evaluated. In this paper, we aim to minimize the gap between training and inference of dense retrievers to achieve better retrieval performance.

However, developing a better dense retriever requires validation, which requires building large indexes from a full corpus (e.g., the entire Wikipedia for open-domain QA) for inference with a huge amount of computational resources and time. To tackle this problem, we first propose an efficient way of validating dense retrievers without building large-scale indexes. Analysis of using a smaller random corpus with different sizes for the validation reveals that the accuracy from small indexes does not necessarily correlate well with the retrieval accuracy on the full index. As an alternative, we construct a compact corpus using a pre-trained dense retriever so that validation on this corpus better correlates well with the retrieval on the full scale while keeping the size of the corpus as small as possible to perform efficient validation.

With our efficient validation, we revisit the training method of dense phrase retrieval (Lee et al., 2021a,b), a general framework for retrieving different granularities of texts such as phrases, passages, and documents. We reduce the training-inference

*JL currently works at Google Research.

†YZ currently works at Apple.

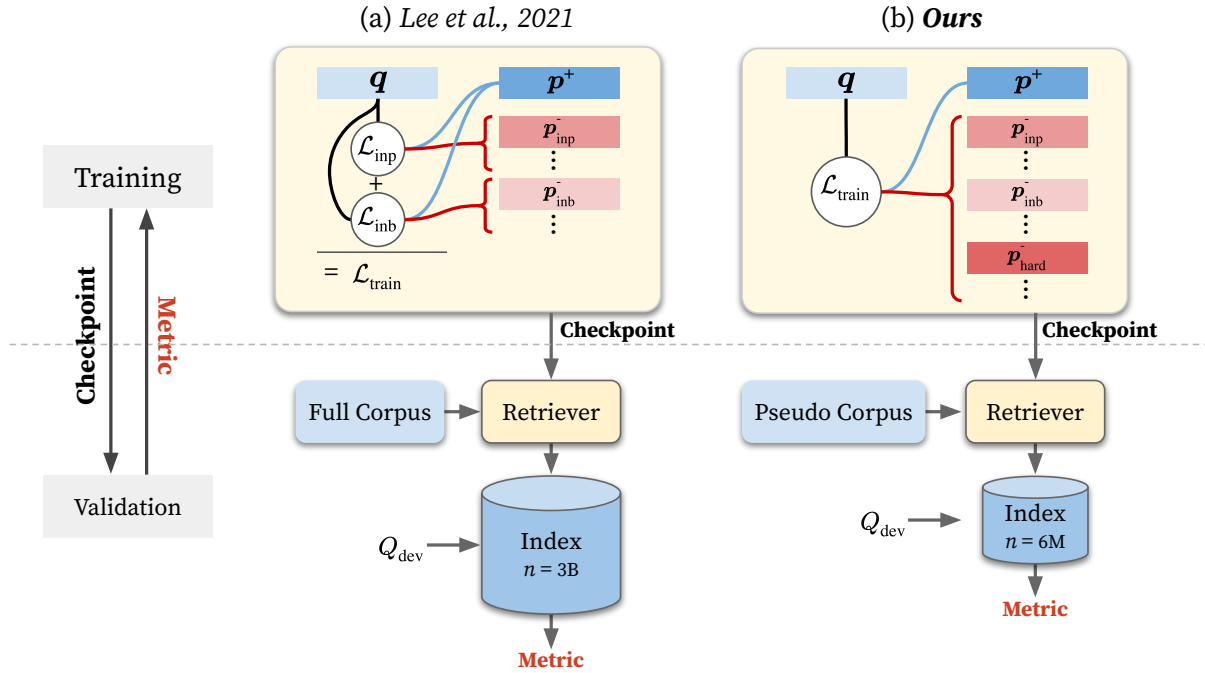


Figure 1: Comparison of the (a) original (Lee et al., 2021a) and (b) proposed procedure for DensePhrases training (top) and validation (bottom). We unify training loss terms \mathcal{L}_{inp} and \mathcal{L}_{inb} that enforce the representation of a question (q) similar to the representation of a positive phrase (p^+) while contrasting from representations of in-passage negative phrases (p_{inp}^-) and in-batch negative phrases (p_{inb}^-) respectively into a single term \mathcal{L}_{train} and expand negatives in number and difficulty with hard negatives (p_{hard}^-). Also, we use a retrieval accuracy on the development set Q_{dev} using a smaller corpus instead of the full corpus as an efficient validation metric for selecting the best checkpoint. Query-side fine-tuning and token filtering are not described in this overview figure.

discrepancy by unifying previous loss terms to discriminate a gold answer phrase from other negative phrases altogether instead of applying in-passage negatives (Lee et al., 2021b) and in-batch negatives separately. To better approximate the retrieval at inference where the number of negatives is extremely large, we use all available negative phrases from training passages to increase the number of negatives and put more weights on negative phrases. We also leverage model-based hard negatives (Xiong et al., 2020) for phrase retrieval, which hasn't been explored in previous studies. This enables our dense retrievers to correct mistakes made at inference time.

Lastly, we study the effect of a representation filter (Seo et al., 2018), an essential component for efficient search. We separate the training and validation of a phrase filtering module to disentangle the effect of contrastive learning and representation filtering. This allows us to do careful validation of the representation filter and achieve a better precision/recall trade-off. Interestingly, we find that a representation filter has a dual role in reducing the index size and also improving retrieval accuracy, meaning smaller indexes are often better than larger

ones in terms of accuracy. This gives a different view of other filtering methods that have been applied in previous studies for efficient open-domain QA (Min et al., 2021; Izacard et al., 2020; Fajcik et al., 2021; Yang and Seo, 2021).

We reemphasize that phrase retrieval is an attractive solution for open-domain question answering compared to other retriever-reader models, considering both accuracy and efficiency. Our contributions are summarized as follows:

- We introduce an efficient method of validating dense retrievers to confirm and accelerate better modeling of dense retrievers.
- Based on our efficient validation, we improve dense phrase retrieval models with modified training objectives and hard negatives.
- Consequently, we achieve the state-of-the-art phrase retrieval accuracy for open-domain QA and also largely improve passage retrieval accuracy on Natural Questions (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017).

2 Related Work

Dense retrieval Retrieving relevant documents for a query (Mitra and Craswell, 2018) is crucial in many NLP applications like open-domain question answering and knowledge-intensive tasks (Petroni et al., 2021). Dense retrievers typically build a search index for all documents by pre-computing the dense representations of documents using an encoder. Off-the-shelf libraries for a maximum inner product search (MIPS) (Johnson et al., 2019; Guo et al., 2020) enable model training and indexing to be developed independently (Lin, 2022). However, both training dense retrievers and building indexes should take into account the final retrieval accuracy. In this respect, we aim to close the gap between training and inference of dense retrievers.

Phrase retrieval Phrase retrieval (Seo et al., 2019) directly finds an answer with MIPS from an index of contextualized phrase vectors. This removes the need to run an expensive reader for open-domain QA. As a result, phrase retrieval allows real-time search tens of times faster than retriever-reader approaches as an alternative for open-domain QA. DensePhrases (Lee et al., 2021a) removes the requirement of sparse features and significantly improves the accuracy from previous phrase retrieval methods (Seo et al., 2019; Lee et al., 2020). Lee et al. (2021b) show how retrieving phrases could be translated into retrieving larger units of text like a sentence, passage, or document, making phrase retrieval a general framework for retrieval. Despite these advantages, phrase retrieval requires building a large index from billions of representations. In this work, we focus on improving phrase retrieval with more efficient validation.

Validation of dense retrieval Careful validation is essential for developing machine learning models to find a better configuration (Melis et al., 2018) or avoid falling to a wrong conclusion. However, many works on dense retrieval do not clearly state the validation strategy, and most of them presumably perform validation on the entire corpus. It is doable but quite expensive¹ to perform frequent validation and comprehensive tuning. Hence, it motivates us to devise efficient validation for dense

¹For example, dense passage retrieval (DPR) (Karpukhin et al., 2020) takes 8.8 hours on 8 GPUs to compute 21-million passage embeddings and 8.5 hours to build a FAISS index. Also, ColBERT (Khattab and Zaharia, 2020) takes 3 hours to index 9M passages in the MS MARCO dataset (Nguyen et al., 2016) using 4 GPUs.

retrieval. Like ours, Hofstätter et al. (2021) construct a small validation set by sampling queries and using a baseline model for approximate dense passage retrieval but limited to early stopping. Liu et al. (2021) demonstrate that small and synthetic benchmarks can recapitulate innovation of question answering models on SQuAD (Rajpurkar et al., 2016) by measuring the concurrence of accuracy between benchmarks. We share the intuition that smaller and well-curated datasets may lead to the same (or sometimes better) model development while faster but with more focus on the validation process.

Hard examples Adversarial data collection by an iterative model (or human) in the loop process aims to evaluate or reinforce models’ weaknesses, including the robustness to adversarial attacks (Kaushik et al., 2021; Bartolo et al., 2021; Nie et al., 2020; Kiela et al., 2021). In this work, we construct a compact corpus from a pre-trained dense retriever for efficient validation. Also, we extract hard negatives from retrieval results of the previous model for better dense representations.

3 Efficient Validation of Phrase Retrieval

Our goal is to train a dense retriever \mathcal{M} that can accurately find a correct answer in the entire corpus C (in our case, Wikipedia). Careful validation is necessary to confirm whether new training methods are truly effective. It also helps find optimal configurations induced by those techniques. However, building a large-scale index for every model makes the model development process slow and also requires huge memory. Thus, an efficient validation method could expedite modeling innovations in the correct direction. It could also allow frequent comparison of different checkpoints when updating a full index simultaneously during the training is computationally infeasible.²

Measuring the retrieval accuracy on an index from a smaller subset of the full corpus (denoted as C^*) for model validation would be a practical choice, hoping $\operatorname{argmax}_{\mathcal{M} \in \Omega} \operatorname{acc}(D|\mathcal{M}, C^*) \approx \operatorname{argmax}_{\mathcal{M} \in \Omega} \operatorname{acc}(D|\mathcal{M}, C)$ where Ω is a set of model candidates and acc means the retrieval accuracy on a QA dataset D . We first examine how a relative order of accuracy between modeling ap-

²Although some works (Guu et al., 2020; Xiong et al., 2020) do asynchronous updates per a specific number of training steps and use the intermediate index for better modeling, it requires a huge amount of computational resources.

proaches may change with varying sizes of the random subcorpus (§3.1) and then develop a clever way to construct a compact subcorpus that maintains reasonable correlation with the retrieval accuracy in the full scale (§3.2).

3.1 Random Subcorpus

Reading comprehension (RC) can be regarded as a special case of open-domain QA, where a corpus contains only a single gold passage (i.e., $C_q = \{c\}$) for each question. Here, the subcorpus is question-dependent. We first gather all gold passages from the development set as a small corpus C_0 , a minimal set that contains answers to all development set questions. We consider a corpus R_r whose size is r times the size of the full corpus by simply appending C_0 with random passages by sampling from the full corpus C , i.e., $C_0 \subset R_r \subset C$ and $|R_r| = r|C|$. We specifically use $r = 1/100, 1/10$ in our experiments. As the corpus size increases, finding the correct answer from a larger number of possible candidates becomes more difficult, so the retrieval accuracy generally decreases (Reimers and Gurevych, 2021).

DensePhrases (Lee et al., 2021a) simply choose the best checkpoint with the highest RC accuracy assuming that a model with better RC accuracy leads to a better retrieval accuracy, or use the last checkpoint at the end of the training.³ It is problematic since our preliminary experiments demonstrate that the RC accuracy and the retrieval accuracy on different sizes of corpus including the full corpus, do not necessarily correlate well with each other. Using a large subcorpus is better for accurate validation not to deviate much from the trends of retrieval accuracy of a full corpus. However, a smaller subcorpus would be better in terms of validation efficiency. This trade-off drives us to design a better way of constructing a validation corpus.

3.2 Hard Subcorpus

The retrieval accuracy given a subcorpus C^* should have a high correlation with the retrieval accuracy over the full corpus, and the size of corpus $|C^*|$ should be small enough (or as small as possible) for efficient validation. For a reasonably accurate dense retriever, it is relatively easy to discriminate a gold phrase from other phrases in random passages. Therefore, it is better to collect a subcorpus with

³RC accuracy generally improves during the training of DensePhrases as the training loss directly optimizes it.

hard passages to test dense retrievers on a similar condition to a full corpus which includes many difficult phrases to discriminate if the corpus can have a limited number of negative passages.

We construct a hard corpus H_k with a compact size using a pre-trained retriever $\bar{\mathcal{M}}$ to extract all context passages of top- k retrieved phrases for all query q in the development set Q_{dev} , and C_0 is merged to always include an answer, i.e., $H_k = C_0 \cup \bigcup_{(q,a) \in Q_{\text{dev}}} \bar{\mathcal{M}}_k(q|C)$ where $\mathcal{M}_k(q|C)$ denotes the top- k passage retrieval results for a query q from the model \mathcal{M} . If $\bar{\mathcal{M}}$ is reasonably accurate, negative examples retrieved by $\bar{\mathcal{M}}$ will make our new model \mathcal{M} difficult to find a correct answer. We expect the retrieval accuracy from H_k quickly drops as k increases and reaches close to the retrieval accuracy on the full corpus C with a manageable k so that we can use retrieval accuracy on a hard subcorpus for efficient validation. It keeps the relative order of models with a much smaller size than the random subcorpus.

4 Optimized Training of DensePhrases

In this section, we briefly review the original training method of DensePhrases (§4.1) and improve it further to reduce the gap between training and retrieval in inference by modifying the training objective (§4.2) and introducing additional training with hard negatives (§4.3).

4.1 Background: Training of DensePhrases

The question encoder and the phrase encoder are jointly trained using reading comprehension datasets. A phrase p is represented as a concatenation of start and end token vectors from the contextualized representations of a context passage c using a phrase encoder. A question q is represented as a concatenation of vectors using two different encoders for the start and the end.

The main training objective is a sum of the two separate contrastive loss terms weighted by the λ coefficient as formally defined in Equation 1.⁴ One is for contrasting a phrase token of positive start/end position (p^+) to that of other positions in the context passage ($N_{\text{inp}} = \{(p; c) \neq (p^+; c) | p \in c\}$). Another is for contrasting the same token to other positive tokens in a current ($N_{\text{inb}} = \{(p'; c') \neq (p^+; c) | (q', p'; c') \in \mathcal{B}\}$) or previous T

⁴We denote the similarity score between a question q and a phrase p as $s(q, p; c)$. While the score and the loss term of start and end tokens are separately calculated in practice, we abbreviate it in the equation for simplicity.

$$\mathcal{L}_{\text{train}}^{\text{org}} = -\log \frac{e^{s(q,p^+;c)}}{e^{s(q,p^+;c)} + \sum_{(p';c) \in N_{\text{inp}}} e^{s(q,p';c)}} - \lambda \log \frac{e^{s(q,p^+;c)}}{e^{s(q,p^+;c)} + \sum_{(p';c') \in N_{\text{inb}} \cup N_{\text{prb}}} e^{s(q,p';c')}} \quad (1)$$

$$\mathcal{L}_{\text{train}} = -\log \frac{e^{s(q,p^+;c)}}{e^{s(q,p^+;c)} + \sum_{(p';c') \in N_{\text{inp}} \cup N_{\text{inb}} \cup N_{\text{prb}} \cup N_{\text{hard}}} \lambda(p^-) e^{s(q,p';c')}} \quad (2)$$

mini-batches ($N_{\text{prb}} = \{(p';c') | (q',p';c') \in \mathcal{B}_{\text{pre}}\}$). The numbers of negatives are $|N_{\text{inp}}| = L - 1$, $|N_{\text{inb}}| = B - 1$, and $|N_{\text{prb}}| = B \times T$ where L is the sequence length of context passages and B ($= |\mathcal{B}|$) is the batch size.

To learn better representations, the dual encoder is first pre-trained with question-answer pairs generated by a question generation model as a data augmentation mainly for better reading comprehension capability and then fine-tuned with original question-answer pairs. Also, knowledge distillation (Hinton et al., 2015) from a stronger reading comprehension model based on a cross encoder to the dual encoder is performed. Lastly, the token filtering classifier (explained more in §5) that discriminates tokens likely to be a start or end of the answer phrases using a linear classifier on top of phrase representations is jointly trained with the dual encoder. We omit two additional loss terms from Equation 1 for knowledge distillation and a token filtering classifier loss for brevity.

Query-side Fine-tuning Documents in the reading comprehension dataset used for the training take only a tiny portion of the entire Wikipedia, and only a small number of negatives for each question-phrase pair are contrasted compared to billion-scale possible phrase candidates in the test time. The query encoder can be further fine-tuned to reduce this discrepancy between training and inference while fixing the phrase encoder and the index by maximizing the likelihood of the gold answer among retrieved phrases for each question. Using more and harder negatives is also an effective way to reduce this gap.

4.2 Unified Loss

The original training objective of DensePhrases (Equation 1) has separate terms for finding a relevant passage (in/pre-batch negatives) and finding the exact phrase position in the passage (in-passage negatives). However, we should find an answer phrase among all possible candidates at once during the test time. Therefore, we modify the loss

term as a unified version (Equation 2) by putting all negatives together into the contrastive targets.

We also introduce the λ coefficient to the unified loss to give weights to negatives. It opens a new question of how we should set the value of λ . The role of λ can be interpreted in two ways. First, multiplying λ to an exponential of a score is equivalent to adding a positive value to the score ($\lambda e^s = e^{s+\log \lambda}$), and then the loss term becomes the soft version of margin-based loss. Second, using λ can mimic the inference time where the number of negative tokens is much larger by duplicating a negative λ times ($\lambda e^s = e^s + e^s + \dots + e^s$) to close the gap between training and test. Based on the second interpretation, we set different value of λ depending on where negative phrase p^- is from: $\lambda(p^-) = \lambda_{\text{inp}} \delta(p^- \in N_{\text{inp}}) + \lambda_{\text{inb}} \delta(p^- \in N_{\text{inb}} \cup N_{\text{prb}}) + \lambda_{\text{hard}} \delta(p^- \in N_{\text{hard}})$.

We extend to use all tokens in context passages with a similar intuition that contrasting with as many tokens as possible could be helpful instead of using only start/end position tokens to in/pre-batch negatives. It changes in/pre-batch negatives to $N_{\text{inb}} = \{(p;c') \neq (p^+;c) | p \in c', (q',p';c') \in \mathcal{B}\}$ and $N_{\text{prb}} = \{(p;c') | p \in c', (q',p';c') \in \mathcal{B}_{\text{pre}}\}$ and their sizes $|N_{\text{inb}}| = B \times L - 1$ and $|N_{\text{prb}}| = B \times T \times L$. This change also increases the number of negatives hundreds of times and turns out empirically advantageous.

4.3 Hard Negatives for Phrase Retrieval

We exploit hard negatives to benefit phrase retrieval, a widely used technique for passage retrieval⁵ but never fully examined for phrase retrieval. We perform a model-based hard negative mining by retrieving top phrases using a pre-trained dual encoder and an index built from this model. We filter

⁵Karpukhin et al. (2020) use one hard negative obtained from BM25 per example in addition to in-batch negatives for training a dual encoder. Xiong et al. (2020) globally select hard negatives from the entire corpus with asynchronously index updates for faster convergence. RocketQA (Qu et al., 2021) denoises hard negatives using cross encoder. The best strategy for hard negative mining and training is still an open problem in dense retrievals.

out phrases whose surrounding passage includes a gold answer (§4.3.1) and then fine-tune the model with extracted hard negatives (§4.3.2). Although we do it only once, this process could be repeated until convergence.

4.3.1 Hard Negative Mining

We use an encoder model and phrase index from the first round to extract model-based hard negatives from top- k phrase retrieval results for questions in the training set. Using high-quality hard negatives by removing false negatives is important to train a better model.

We exclude examples when a context passage of a retrieved phrase contains an answer. A context passage corresponding to a retrieved phrase can be restored using information stored along with the index. It helps to focus more on topically different documents and shares the intuition from the analysis in Lee et al. (2021b) that DensePhrases rely less on topical relevance than DPR. Compared to a more strict condition based on the exact match that may miss almost correct phrases with a minor error in the boundary by misclassifying the exact position, it reduces about 20% of negative pairs, hopefully reducing false negatives and achieving higher accuracy gain. Besides, these rules are somewhat loose in that there could be multiple possible answers to a question, and different representations for the same entity could exist since the annotated answer list is imperfect. We left filtering based on a cross encoder (Qu et al., 2021; Ren et al., 2021) to future work due to the convenience of automatic filtering.

4.3.2 Training with Hard Negatives

After the hard negative mining, we fine-tune a dual encoder with the hard negatives. We sample h hard negatives for each training step and append them to negative targets for the loss calculation.⁶ We expect that hard negatives give a better training signal than random in/pre-batch negatives (Xiong et al., 2020) because those are examples difficult to discriminate for the previous model. Moreover, hard negatives extracted from the larger corpus could expose a model to other diverse documents than the original training dataset. This is similar to query-side fine-tuning but differs in that both the question encoder and phrase encoder are updated.

⁶If the number of hard negatives after removing false negatives is less than h , we sample random passages to match the number of hard negatives for parallel computation.

There are different possible options for choosing N_{hard} . First, we may include only corresponding hard negatives for each example or all hard negatives in a mini-batch. Second, we may include only each negative’s start/end position token or all tokens in the context passage. Similar to §4.2, we include all tokens in the context passage of all hard negatives in a mini-batch for N_{hard} . Using all available tokens is generally better because they potentially belong to the final negative phrase candidates in inference. Training with larger numbers of negatives is beneficial to reduce the gap between training and inference. Including all of them does not induce significant additional memory overhead since we should encode the same number (h) of passages regardless of different options. Therefore, we use $N_{\text{hard}} = \bigcup_{(\hat{p}; \hat{c}) \in H(q, p^+; c), (q, p^+; c) \in \mathcal{B}} \{(p; \hat{c}) | p \in \hat{c}\}$ as all tokens from all hard negatives in a mini-batch where H is a set of the sampled h hard negatives and $|N_{\text{hard}}| = B \times L \times h$.

5 Token Filtering

Representation filtering is often applied in practice to reduce the index size for efficient search (Min et al., 2021). For phrase retrieval, tokens that are not likely to be a start/end position of an answer are filtered out using a trained filter classifier based on a logit score for each token to reduce an index size without losing much accuracy. Only tokens with a score larger than a specific threshold are kept. After the filtering, the index is compressed using optimized product quantization (Ge et al., 2013).

5.1 Token Filter Threshold

A filter threshold for the token filter determines a trade-off between the index size (efficiency) and retrieval accuracy (Figure 2). Interestingly, we find that token filtering can even improve retrieval accuracy. As we increase a threshold from a very small value (not filtering), the accuracy fluctuates but generally increases until a specific threshold because the filter successfully reduces the number of candidates, making prediction easier. After that threshold, the accuracy drops quickly because the filter starts to leave out many correct tokens.

However, finding the peak retrieval accuracy requires a manual search of different thresholds after indexing and evaluating. Since using it as a validation metric is expensive, we first select the best checkpoint based on retrieval accuracy without performing any token filtering. Especially when




Figure 2: The trade-off between the index size and validation retrieval accuracy by changing filter thresholds on random subcorpora with different sizes, (a) C_0 , (b) $R_{1/100}$, and (c) $R_{1/10}$. A threshold that gives better accuracy with a smaller index size exists. Acc@1 (blue) is more unstable than Acc@10 (red). Interestingly, the index size of peak EM@1 is smaller than that of peak Acc@10.

the token filter is in the middle of training, the peak threshold will vary, and using a specific fixed threshold would not be fair. Also, the best threshold changes depending on the corpus size, so choosing a threshold for the full corpus based on a smaller corpus is not straightforward.

5.2 Token Filter Training and Valiation

Lee et al. (2021a) jointly train a token filter classifier with a dual encoder. It is convenient in that an additional training process is not required, while we should tune the weight for a loss before adding to the overall training loss. Training pushes phrase vectors into two moving cones toward the start and end vectors since a logit is a dot product score between a phrase vector and a start/end vector. It has two potential disadvantages: (1) phrase representations are concentrated on a subset of the entire feature space, so the expressiveness of the model is not fully exploited, and (2) optimization is more difficult because of the moving targets.

To address the issues, we train a token filter after training a phrase encoder. We expect that the two-stage training process encourages phrase representations to be distributed over the space. Moreover, we can validate the token filter separately due to the separate training process and pick the best one. We cannot decide the threshold during the filter training, so we use the AUC-PR metric for filter validation by measuring precision and recall by sweeping all threshold values.

6 Experiments

To show the effectiveness of our proposed method, we evaluate DensePhrases models on open-domain QA benchmarks following the experimental setup

of Lee et al. (2021a,b).⁷

Datasets We measure phrase retrieval accuracy and passage retrieval accuracy on two open-domain QA datasets following the standard train/dev/test splits: Natural Questions (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017). We first train our phrase retrieval models on Natural Questions (DensePhrases[♥]) or on multiple reading comprehension datasets (DensePhrases[♠]), namely Natural Questions, WebQuestions (Berant et al., 2013), CuratedTREC (Baudiš and Šedivý, 2015), TriviaQA, and SQuAD (Rajpurkar et al., 2016). Then each model is further query-side fine-tuned on Natural Questions and TriviaQA. We build the phrase index with smaller subsets of corpora (R_r or H_k) for validation and use the 2018-12-20 Wikipedia snapshot (C) for the final inference.

Training details We use the same training hyperparameters of the original DensePhrases except for the batch size $B = 48$. We set the number of training epochs to 2 with the generated question-answer pairs and the number of training epochs to 10 with the standard reading comprehension dataset for more careful validation. We set $\lambda_{\text{inb}} = 256$ and $\lambda_{\text{inp}} = \lambda_{\text{hard}} = 1$. We set $k = 10$ and $h = 1$ for the hard negative mining and sampling.

Token filtering Our token filter achieves an improved AUC-PR value over the filter from the original DensePhrases model (e.g., 0.348 vs. 0.307). We use a filter threshold of -3.0 for the index with the full corpus. This threshold reduces the index size by more than 70% of the original size.

⁷<https://github.com/princeton-nlp/DensePhrases>

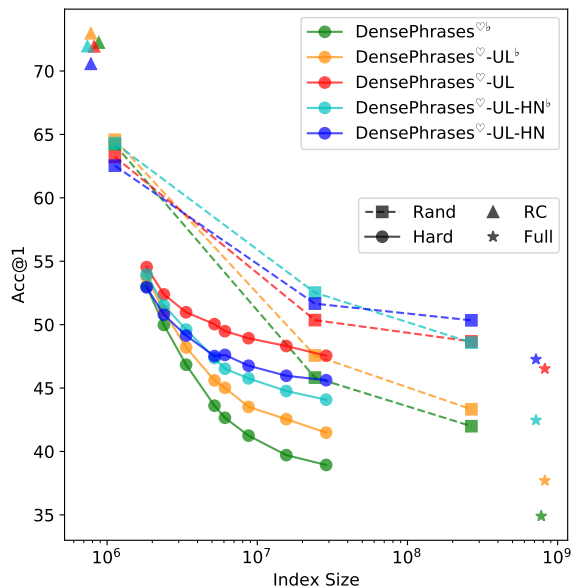


Figure 3: Validation results on open-domain QA. We plot retrieval accuracy (Acc@1) on indexes with different sizes (log-scale) from random and hard subcorpora. Random subcorpora (■) starts with C_0 and is extended to $R_{1/100}$ and $R_{1/10}$. Hard subcorpora (●) include H_k for $k \in \{1, 2, 4, 8, 10, 16, 32, 64\}$. We also plot reading comprehension (RC) accuracy⁸ and retrieval accuracy on the full index with filtering. We compare five different models with or without our proposed training methods and query-side fine-tuning. All models are trained and evaluated on Natural Questions. UL, HN, and ^b indicate a model trained with the unified loss, hard negatives, and before query-side fine-tuning.

6.1 Model Validation

In our preliminary experiments, we observe that the best checkpoint among training epochs differs depending on the corpus size (especially for small scale). Figure 3 shows the validation retrieval accuracy of the DensePhrases models with different training methods on various sizes of random and hard subcorpora. The retrieval accuracy on the hard subcorpus rapidly drops and reaches close to the retrieval accuracy on the full corpus as k increases with moderately increasing index size. On the other hand, retrieval accuracy on a random subcorpus is higher than on a hard subcorpus with a similar index size. For instance, retrieval accuracies on H_8 (5.1M) are lower than those on $R_{1/100}$ (24.2M) with 4 times smaller index, and retrieval accuracies on H_{16} (8.7M) are lower than those on $R_{1/10}$ (266.4M) with 30 times smaller index. It indicates

⁸Since the length of a passage for each question varies, we put aside points corresponding to RC on the left of the figure with arbitrary small index sizes.

Model	NQ	TQA
	Acc@1	Acc@1
DPR [◇] + BERT reader	41.5	56.8
DPR [♣] + BERT reader	41.5	56.8
RePAQ [◇] (retrieval-only)	41.2	38.8
RePAQ [♣] (retrieval-only)	41.7	41.3
DensePhrases [♡]	40.9	50.7
DensePhrases [♣]	41.3	53.5
DensePhrases [♡] -UL	43.5	51.3
DensePhrases [♡] -UL-HN	44.0	47.0
DensePhrases [♣] -UL	42.4	55.5

Table 1: Open-domain QA phrase retrieval test results. We report top-1 accuracy (Acc@1). [◇]: trained on each dataset independently. [♣]: trained on multiple datasets. [♡]: trained on Natural Questions datasets.

that a hard subcorpus can effectively imitate inference with a full corpus, where correct retrieval is the most difficult.

The relative order of accuracy between models on hard subcorpus converges quickly at around H_{10} (6.1M). However, the order of accuracy when using random subcorpus changes from $R_{1/100}$ to $R_{1/10}$ showing the difficulty of efficient validation. On the other hand, retrieval accuracy on a hard subcorpus is more stable and serves as an efficient validation metric.

Our validation results clearly demonstrate that unified loss is helpful. Query-side fine-tuning also harms the RC accuracy and the retrieval accuracy with C_0 (1.1M) while improving the retrieval accuracy with larger indexes. It shows how a wrong conclusion can be made from small-sized corpora.

6.2 Phrase Retrieval

Table 1 summarizes end-to-end open-domain QA results. Both unified loss and hard negatives are shown to be effective. With our improved training methods, the best model surpasses the original DensePhrases model by 2.7 points in Natural Questions and 2.0 points in TriviaQA, achieving the state-of-the-art retrieval-only open-domain QA performance.

6.3 Passage Retrieval

Table 2 summarizes open-domain QA passage retrieval results. Our method also improves passage retrieval accuracy significantly. The best model improves top-20 passage retrieval accuracy by 4.0 points in Natural Questions and 1.8 points in TriviaQA. It again shows that DensePhrases can be used for passage retrieval as well. We may

Model	Natural Questions					TriviaQA				
	Acc@1	Acc@5	Acc@20	MRR@20	P@20	Acc@1	Acc@5	Acc@20	MRR@20	P@20
DPR [◇]	46.0	68.1	79.8	55.7	16.5	54.4	-	79.4	-	-
DPR [♠]	44.2	66.8	79.2	54.2	17.7	54.6	70.8	79.5	61.7	30.3
DensePhrases [♡]	50.1	69.5	79.8	58.7	20.5	-	-	-	-	-
DensePhrases [♠]	51.1	69.9	78.7	59.3	22.7	62.7	75.0	80.9	68.2	38.4
DensePhrases [♡] -UL	57.1	75.7	83.7	65.2	22.0	62.0	74.6	80.6	67.6	33.3
DensePhrases [♡] -UL-HN	58.6	75.7	83.4	66.1	21.9	60.3	73.3	79.6	66.1	32.3
DensePhrases [♠] -UL	56.7	75.9	83.8	65.2	23.7	65.0	76.6	82.7	70.2	39.0

Table 2: Open-domain QA passage retrieval test results. We report top- k passage retrieval accuracy (Acc@ k , for $k \in \{1, 5, 20\}$), mean reciprocal rank at 20 (MRR@20), and precision at 20 (P@20). [◇]: trained on each dataset independently. [♠]: trained on multiple datasets. [♡]: trained on Natural Questions datasets.

use DensePhrases as a building block of other tasks and expect to achieve good phrase retrieval performance with expressive reader models like FiD (Izacard and Grave, 2021).

6.4 Discussion on Hard Negatives

From Figure 3, we observe that with hard sub-corpora, the model trained with hard negatives (cyan) shows higher validation accuracy than the model without hard negative training (yellow) before query-side fine-tuning, but their order changes after query-side fine-tuning (blue vs. red). This is because the hard negative mining process is similar to hard corpus construction, blurring the accurate estimation of validation performance. However, we pick the best model before the query-side fine-tuning, which lets us decide to go with hard negatives (due to cyan vs. yellow) and achieve state-of-the-art performance with the full index.

From Table 1, we observe that hard negatives improve in-domain accuracy but harm the out-of-domain accuracy. Since hard negative passages are close to the original training data, it improves the performance of questions from the same domain but could cause overfitting and harm the generalization ability. This observation solicits better hard negative mining methods.

7 Conclusion

In this study, we aim to bridge the gap between training and inference of phrase retrieval. We first develop an efficient validation metric that measures retrieval accuracy on the index from a small corpus with hard passages using a pre-trained retriever. Based on this validation, we show that the improvements in dense phrase retrieval training with unified loss and hard negatives are effective. As a result, we achieve state-of-the-art phrase retrieval and pas-

sage retrieval accuracy in open-domain question answering among retrieval-only approaches.

Our work demonstrates that thorough validation is crucial for the accurate and efficient development of phrase retrieval with a large corpus. Also, we prove that modeling and training methods should be designed closely to retrieval in inference time. Despite its remarkable efficiency and flexibility, phrase retrieval has been relatively less studied than passage retrieval. We believe that our work can encourage more study on phrase retrieval with an efficient development cycle. Furthermore, we hope our findings could be extended to dense retrieval in general to help a wide variety of applications. Moreover, it could be especially beneficial in real applications where the corpus size is much larger than benchmark datasets.

Limitations

This work focuses on phrase retrieval, where the training-inference discrepancy might be more significant than other dense retrieval cases, based on the DensePhrases (Lee et al., 2021a) framework. We plan to explore other dense retrieval methods in the future. We use open-domain question answering as the main benchmark to show the effectiveness of the proposed method but expect a wide application to other knowledge-intensive tasks.

Acknowledgements

The authors appreciate Alon Albalak, Sungdong Kim, Sewon Min, Wenhao Yu, and anonymous reviewers for proofreading the manuscript and giving constructive feedback. The research presented in this work was funded by Meta AI. The views expressed are those of the authors and do not reflect the official policy or position of the funding agency.

References

- Max Bartolo, Tristan Thrush, Robin Jia, Sebastian Riedel, Pontus Stenetorp, and Douwe Kiela. 2021. [Improving Question Answering Model Robustness with Synthetic Adversarial Data Generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8830–8848.
- Petr Baudiš and Jan Šedivý. 2015. [Modeling of the Question Answering Task in the YodaQA System](#). In *Proceedings of the 6th International Conference on Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 222–228. Springer.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic Parsing on Freebase from Question-Answer Pairs](#). In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to Answer Open-Domain Questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 539–546. IEEE.
- Martin Fajcik, Martin Docekal, Karel Ondrej, and Pavel Smrz. 2021. [Pruning the Index Contents for Memory Efficient Open-Domain QA](#). *arXiv preprint arXiv:2102.10697*.
- Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. [Optimized Product Quantization](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(4):744–755.
- Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. [Accelerating Large-Scale Inference with Anisotropic Vector Quantization](#). In *International Conference on Machine Learning*, pages 3887–3896. PMLR.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: Retrieval-Augmented Language Model Pre-Training](#). In *International Conference on Machine Learning*.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. [Distilling the Knowledge in a Neural Network](#). *arXiv preprint arXiv:1503.02531*.
- Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. [Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 113–122.
- Gautier Izacard and Édouard Grave. 2021. [Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880.
- Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola De Cao, Sebastian Riedel, and Edouard Grave. 2020. [A Memory Efficient Baseline for Open Domain Question Answering](#). *arXiv preprint arXiv:2012.15156*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. [Billion-Scale Similarity Search with GPUs](#). *IEEE Transactions on Big Data*, 7(3):535–547.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781.
- Divyansh Kaushik, Douwe Kiela, Zachary C Lipton, and Wen-tau Yih. 2021. [On the Efficacy of Adversarial Data Collection for Question Answering: Results from a Large-Scale Randomized Study](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6618–6633.
- Omar Khattab and Matei Zaharia. 2020. [ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, et al. 2021. [Dynabench: Rethinking Benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. [Natural Questions: A Benchmark for Question Answering Research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.

- Jinhyuk Lee, Minjoon Seo, Hannaneh Hajishirzi, and Jaewoo Kang. 2020. [Contextualized Sparse Representations for Real-Time Open-Domain Question Answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 912–919.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021a. [Learning Dense Representations of Phrases at Scale](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6634–6647.
- Jinhyuk Lee, Alexander Wettig, and Danqi Chen. 2021b. [Phrase Retrieval Learns Passage Retrieval, Too](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3661–3672.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent Retrieval for Weakly Supervised Open Domain Question Answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. [PAQ: 65 Million Probably-Asked Questions and What You Can Do With Them](#). *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Jimmy Lin. 2022. [A Proposed Conceptual Framework for a Representational Approach to Information Retrieval](#). In *ACM SIGIR Forum*, volume 55-2, pages 1–29. ACM New York, NY, USA.
- Nelson F Liu, Tony Lee, Robin Jia, and Percy Liang. 2021. [Can Small and Synthetic Benchmarks Drive Modeling Innovation? A Retrospective Study of Question Answering Modeling Approaches](#). *arXiv preprint arXiv:2102.01065*.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. [On the State of the Art of Evaluation in Neural Language Models](#). In *International Conference on Learning Representations*.
- Sewon Min, Jordan Boyd-Graber, Chris Alberti, Danqi Chen, Eunsol Choi, Michael Collins, Kelvin Guu, Hannaneh Hajishirzi, Kenton Lee, Jennimaria Palomaki, et al. 2021. [NeurIPS 2020 EfficientQA competition: Systems, analyses and lessons learned](#). In *NeurIPS 2020 Competition and Demonstration Track*, pages 86–111. PMLR.
- Bhaskar Mitra and Nick Craswell. 2018. [An Introduction to Neural Information Retrieval](#). Now Foundations and Trends.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A human generated machine reading comprehension dataset](#). In *CoCo@ NIPS*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A New Benchmark for Natural Language Understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, et al. 2021. [KILT: a Benchmark for Knowledge Intensive Language Tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Nils Reimers and Iryna Gurevych. 2021. [The Curse of Dense Low-Dimensional Information Retrieval for Large Index Sizes](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 605–611.
- Ruiyang Ren, Yingqi Qu, Jing Liu, Wayne Xin Zhao, Qiaoqiao She, Hua Wu, Haifeng Wang, and Ji-Rong Wen. 2021. [RocketQAv2: A Joint Training Method for Dense Passage Retrieval and Passage Re-ranking](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2825–2835.
- Minjoon Seo, Tom Kwiatkowski, Ankur Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2018. [Phrase-Indexed Question Answering: A New Challenge for Scalable Document Comprehension](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 559–564.
- Minjoon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. [Real-Time Open-Domain Question Answering with Dense-Sparse Phrase Index](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4430–4441.
- Aaron Van den Oord, Yazhe Li, and Oriol Vinyals. 2018. [Representation Learning with Contrastive Predictive Coding](#). *arXiv e-prints*, pages arXiv–1807.

Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. [Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval](#). In *International Conference on Learning Representations*.

Sohee Yang and Minjoon Seo. 2021. [Designing a Minimal Retrieve-and-Read System for Open-Domain Question Answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5856–5865.